REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY): 04-03-2010

2. REPORT TYPE: Final Report

3. DATES COVERED (From - To): 1-Sep-2002 - 31-Aug-2008

4. TITLE AND SUBTITLE: Final Report: Adaptive Coordinated Control of Intelligent Multi-Agent Teams

5a. CONTRACT NUMBER: DAAD19-02-1-0383

5b. GRANT NUMBER:

5c. PROGRAM ELEMENT NUMBER: 611103

5d. PROJECT NUMBER:

5e. TASK NUMBER:

5f. WORK UNIT NUMBER:

6. AUTHORS

S. Shankar Sastry, R. Vijay Kumar, Howard M. Choset

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES

University of California - Berkeley
Sponsored Projects Office
The Regents of the University of California
Berkeley, CA 94704 -5940

8. PERFORMING ORGANIZATION REPORT NUMBER:

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S): ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S): 43550-NS-MUR.59

12. DISTRIBUTION AVAILABILITY STATEMENT

Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

The goal of this MURI program (ACCLIMATE) has been to develop adaptive, intelligent, and multi-agent control technologies to address the increasing concerns on the reliability, adaptability, and robustness of complex systems (such as unmanned air/ground vehicles) encountered on modern battlefields or in homeland defense. In particular, the program has been focused on the following three thrusts:
1. Architecture design of adaptive heterogeneous systems. We design novel system architectures that organically

15. SUBJECT TERMS

multi-agent control, adaptive control, hybrid systems, information fusion

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
UU	UU	UU	UU		S. Shankar Sastry
					19b. TELEPHONE NUMBER: 510-642-5771

Final Report

Adaptive Coordinated Control of Intelligent Multi-Agent Teams

ARO MURI DAAD19-02-1-0383

S. Shankar Sastry (PI)

Department of Electrical Engineering and Computer Sciences
514 Cory Hall, Berkeley CA 94720-1770
(510) 642-5771 (office) (510) 643-8426 (fax)
sastry@coe.berkeley.edu

TEAM MEMBERS:

UC Berkeley: S. Sastry (PI), L. El Ghaoui, M. Jordan, J. Malik, S. Russell, and P. Varaiya
UPenn: V. Kumar (co-PI), R. Bajcsy, K. Daniilidis, J. Ostrowski, G. Pappas, and C. Taylor
CMU: H. Choset (co-PI), A. Rizzi, and C. Thorpe

The goal of this MURI program (ACCLIMATE) has been to develop adaptive, intelligent, and multi-agent control technologies to address the increasing concerns on the reliability, adaptability, and robustness of complex systems (such as unmanned air/ground vehicles) encountered on modern battlefields or in homeland defense. In particular, the program has been focused on the following three thrusts:

1. Architecture design of adaptive heterogeneous systems. We design novel system architectures that organically incorporate human intervention, and can be composed interoperably on the fly.

2. Integration of multi-sensor information into 3-D virtual environments for incorporating human intervention.

3. Handling uncertainty and adversarial intent in adaptive planning.

In this final report, we summarize and highlight the key contributions of this program towards achieving the above fundamental goals.

1 Thrust I: Architecture Design of Adaptive Heterogeneous Systems.

1. A flight control system for aerial robots.

   We have proposed a hierarchical flight control system for unmanned aerial vehicles [4]. The proposed system executes high-level mission objectives by progressively substantiating them into machine-level commands. The acquired information from various sensors is propagated back to the higher layers for reactive decision making. Each vehicle is connected via standardized wireless communication protocol for scalable multi-agent coordination.

   Figure 1 shows the overall layers in the hierarchical flight management system. First, motion-related information that is vital for vehicle control and high-level operation is measured by the onboard navigation sensors in the *sensing* layer. These mobile sensors include inertial navigation system (INS), global positioning system (GPS), ultrasonic sensors, and laser trange-finders. Second, the *reasoning and coordination* layer follows. Each component in this layer is responsible for strategy planning of a specific task. Any appropriate strategy planner for a given mission is then selected by a switching layer. Finally, the *action* layer instructs the UAV to move to strategic locations that are computed by the decision making process.

   The proposed system has been successfully implemented on a number of small helicopters and validated in various applications, including *waypoint navigation*, *pursuit-evasion game*, and *target tracking*.

2. Pursuit-evasion strategies for teams of multiple agents with incomplete information.

   We have investigated search strategies for multi-player pursuit-evasion games, in which a team of pursuers try to detect and capture multiple intelligent evaders [1]. These are games with incomplete information for which there is no complete theory of existence of optimal solutions. In the pursuit-evasion game that we consider, a pursuer (or a group of pursuers) searches for an evader (or a group of evaders). A suitable physical analogy would be a military surveillance mission of a public safety operation with unmanned aerial vehicles (UAVs) and unmanned ground vehicles (UGVs) serving as the pursuers and evaders, respectively. In our setting, all the players act in a bounded rectangular area. The pursuers can sense only a finite region around them, which presents the difficulty of partial observability. Based on the history of such limited sensing information, the pursuers try outsmart the evaders and eventually capture them.

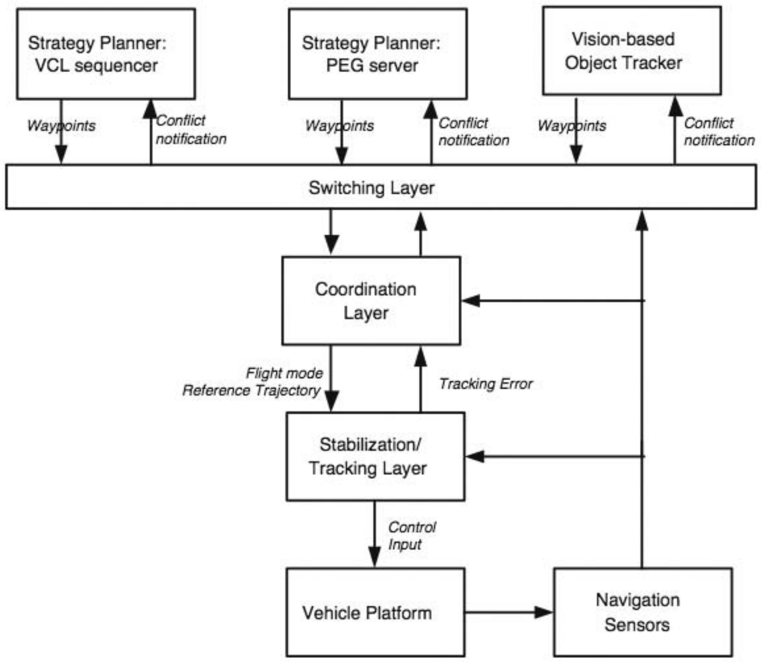

Figure 1: The multi-functional hierarchical flight management system.

In this work, we first incorporated overlap-reducing algorithms to enable multiple pursuers to chase multiple evaders. These new algorithms reduce the amount of sensing overlap between pursuers, thus distributing them more efficiently. Furthermore, we developed heuristic algorithms and policies for pursuers with varying sensing capabilities against faster evaders with enhanced intelligence.

The effectiveness of the proposed pursuit policies was evaluated via simulation. The results suggested that the *globa-subregional* and the *local-max* policies perform the best among the five tested pursuit policies. The overall better performance of these two policies can be attributed to two factors: 1. The reduce sensing overlap between pursuers. 2. They do not over-assign pursuers to target locations. Our analysis also provides a basis for developing policy classes for the study of finite horizon games.

3. Synchronization in oscillator networks.

   Understanding how synchronization is achieved in a network provides us with invaluable information for constructing large-scale systems. For example, in leaderless coordination of multi-agent systems, a key issue is how to achieve consensus for arbitrary topologies. Perhaps the most celebrated model for synchronization is the Kuramoto model, a system of structured ordinary differential equations, which has been used to explain how synchronization is achieved in many engineering, physical, and biological systems.

   In our work [11], we have considered the problem of synchronization of coupled oscillator networks in the Kuramoto framework with heterogeneous communication delays. We have manifested that synchronization can be achieved even if the oscillators are non-identical, if we allow heterogeneous delays in the networks. Moreover, we have established synchronization for Kuramoto network with switching topologies even in the presence of delays.

   We have shown that the synchronization of the delayed Kuramoto models may not be global in the nonlinear case, as in general the solution of the consistency may not be unique. A nonlinear result similar to the one obtained in the case of network congestion control would be desirable, at least in the case in which there is a unique attracting set. This will be the focus of our future research.

4. On consensus over random networks.

   Over the past few years, algorithms for solving different variants of the decentralized consensus problem have received a significant amount of attention. In general, the main goal of such algorithms is to achieve a global objective over a network of agents using purely local interactions. Due to absence of a centralized computational entity, possible lack of information on global topology and limited energy resources, one desirable property of such algorithms is robustness to occurrence of discontinuous changes in the topology of the multiagent network. These topology changes may occur naturally due to communication constraints and/or movement of agents.

   Our project [12] presented a more general framework, compared to the already existing results, for almost sure convergence of the agreement protocols for a randomly changing network. Our framework contains both the random gossip algorithm with synchronous time model and the Erdos-Renyi random graph model as special cases. Moreover, the results are more general comparing to those that only provide a sufficient condition for reaching consensus in probability, while our theorems provide a necessary and sufficient condition for almost sure convergence.

   The main assumption behind our model is that the underlying random graphs of the network are mutually independent over time and have the same distribution. We also assume that the parameters of the distributed consensus algorithm at any given time step are independent of the parameters at all other times and are drawn from the same probability distribution. In the future, we plan to extend our results for the more general case of dependence over time to include the case of agents with motion.

2 Thrust II: Integration of Multi-Sensor Information into 3-D Virtual Environments for Incorporating Human Intervention.

1. Wide area multiple camera calibration.

   The calibration of cameras distributed in a wide area is a challenging task because it is impossible to use reference objects visible to all cameras and because wide field-of-view cameras suffer under *radial distortion*. Our work [2] proposes the first algorithm in the literature for radial distortion estimation from multiple views without involving non-linear minimization. The correspondences between views are obtained by deliberately moving an LED in thousands of unknown positions in front of the cameras. Then both projection matrices and radial distortion parameters are simultaneously computed using a factorization approach.

   The main contribution of this paper is that the simultaneous treatment of projection matrices and radial distortion over multiple views and without any nonlinear minimization. The algorithm differs from others in the literature for being a one-shot method with minimum computational effort. Moreover, it makes use of all multiple-view constraints simultaneously, instead of just pairwise fundamental matrices. Compared with the existing methods, it provides an estimation for the projection matrices and radial distortion without requiring Euclidean stratification and nonlinear minimization.

2. Experiments on visual loop closing using vocabulary trees.

   Imagine a robot vehicle traveling in an urban environment with a camera. The goal of visual loop closing is to detect when the robot is close to a previously visited location based on the images acquired previously. Since we do not assume that any image of the environment is previously available before the current acquisition, the whole process must be performed in an online fashion. Hence, to attain reasonable computational time, it is imperative to avoid comparing the current image to all previously seen images. In our experiments, we used GPS positioning as a prior.

In this work [5], we studied the problem of visual loop closing for long trajectories in an urban environment. Our approach is based on the vocabulary tree method, which is built to define a set of visual words. Features are said to be similar if they correspond to the same visual word. The procedure can be summarized as follows. A vocabulary tree is built offline. Image descriptors from each image of the sequence are added sequentially to the vocabulary tree using inverted files. When the vehicle re-approaches a previously visited location, the descriptors of the current image are used in conjunction with the inverted files to obtain the N closest images. Then, geometric consistency using the epipolar constraint is used to determine if the matches are good. We experiment with two types of trees and with features from the trajectory as well as with features completely independent of the trajectory.

3. Digitizing archaeological excavations from multiple views.

   We proposed a novel approach on digitizing large scale unstructured environments like archaeological excavations using off-the-shelf digital still cameras [14]. The cameras are calibrated with respect to few markers captured by a theodolite system. Our algorithm has as input multiple calibrated images and outputs an occupancy voxel space where occupied pixels have a local orientation and a confidence value. Both orientation and confidence facilitate an efficient rendering and texture mapping of the resulting point cloud.

   The algorithm combines the following new features: Images are backprojected to hypothesized local patches in the world and correlated on these patches yielding the best orientation. Adjacent cameras build tuples which yield a product of pairwise correlations, called strength. Multiple camera tuples compete each other for the best strength for each voxel. A voxel is regarded as occupied if strength is maximum along the normal. Unlike other multi-camera algorithms using silhouettes, photoconsistency, or global correspondence, our algorithm makes no assumption on camera locations being outside the convex hull of the scene. Figure 2 demonstrates a reconstruction sequence.

4. Correspondence-free structure-from-motion.

   Estimation of 3-D motion from two calibrated views has been exhaustively studied in the case where optical flow or feature correspondences are given and the scene is rigid. Algorithms working over multiple frames yield high-quality motion trajectories and reconstructions where feature matches are cleaned through outlier rejection and motions independent of the camera are excluded. These outlier rejection and segmentation steps are subject to fundamental coupling of data association and estimation: If we knew the motion estimate, data association would be trivial; if we knew the data association, motion estimation would be easier. Resistance to outliers and independent motions pose severe practical limitations to the wide application of structure from motion as a navigation tool, visual GPS, or a camera tracker.

   we have proposed a novel approach for structure-from-motion applicable in the presence of large motions and many irrelevant features resulting from reduced overlap of the fields of view [9]. The approach is based on the naive principle that an exhaustive search over all possible correspondence configurations for all motion hypotheses would yield all 3-D motions compatible with these two views. Such a search is nevertheless intractable when we use a large field of view. The contribution of this paper is in the reformulation of this Hough-reminiscent approach as a filtering problem: Assuming a similarity function between any two features in the first and second view, we convolve this function with a kernel that checks the compatibility of a correspondence pair with the epipolar constraint

(a) Two original image pairs.

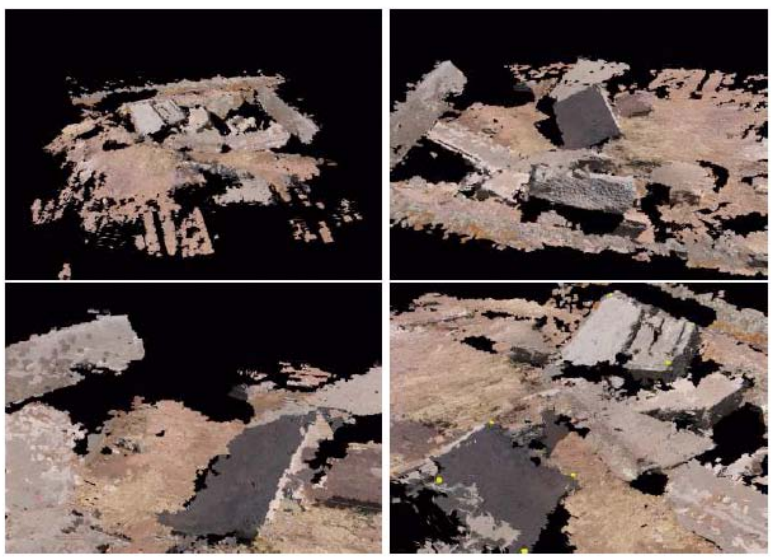

(b) Reconstruction of the scene.

Figure 2: Illustration of a reconstruction sequence.

for a given motion hypothesis. The resulting integral is a *Radon* transform known from computer tomography where a material density is integrated over a ray path.

We present a complete end-to-end system, from images to motion parameters where the only tuning parameter is the coupled resolution of the image and the motion space. We extract SIFT features for which we define their similarity function proportional to the Euclidean norm of the attribute vectors and we compute the spherical harmonics of the similarity function as the input to the correlation integral. The results on real sequences compare favorably to traditional robust estimation methods.

5. Full automatic registration of 3-D point clouds. During the last few years we have experienced the market introduction of range sensors with reasonable cost as well as the availability of many successful

stereo vision algorithms. As any registration problem, range registration consists of the steps of matching and estimation of the rigid transformation. Depending on the displacement and orientation between point clouds, we differentiate between crude and fine alignment. The challenge in crude registration lies in performing it automatically and consistently even when there is very small overlap. The best existing technique for fine registration is the Iterative Closest Point (ICP) algorithm.

We have proposed a novel technique for the registration of 3-D point clouds which makes very few assumptions [8]. We avoid any manual rough alignment or the use of landmarks, displacement can be arbitrarily large, and the two point sets can have very little overlap. Crude alignment is achieved by estimation of the 3-D rotation from two Extended Gaussian Images (EGIs) even when the data sets inducing them have partial overlap. The technique is based on the correlation of the two EGIs in the Fourier domain and makes use of the spherical and rotational harmonic transforms. For pairs with low overlap which fail a critical verification step, the rotational alignment can be obtained by the alignment of constellation images generated from the EGIs. Rotationally aligned sets are matched by correlation using the Fourier transform of volumetric functions. A fine alignment is acquired in the final step by running Iterative Closest Points with just a few iterations.

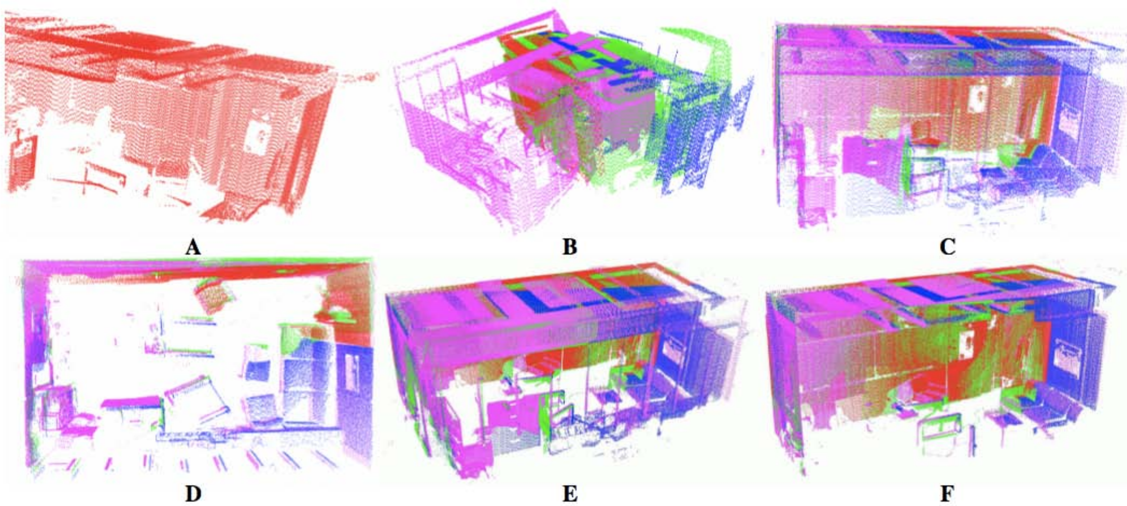

Figure 3: (A) A representative room range scan. (B) Poor alignment obtained by running ICP on the input. (C,D) A side and overhead view of the rough alignment. (E,F) A full and partial view of the final alignment.

6. Single-viewpoint, catadioptric cone mirror omnidirectional imaging theory and analysis.

   Ordinary cameras used in machine vision either have a narrow field of view (FOV) or have a wide FOV but suffer from complex distortion. It can be difficult to unwarp a wide FOV image to obtain perspective projection views accurately. Based purely on the ideal projection imaging model, it has been shown that surfaces of revolution of conic section curves are the only mirror shapes that can be paired with a single converging projection camera to create single-viewpoint (SVP), catadioptric omnidirectional view systems whose omniview image can be unwarped to perspective projection views without systematic distortions. The pinhole-model-based geometry has also been analyzed by others. The condition critical to being able to unwarp to perspective projection views from a single omniview image is the SVP condition. The cone shape, although a surface of revolution of a conic section, has been deemed unusable because, it had been said, in its SVP configuration only rays grazing the mirror surface can be seen. In fact, practical SVP cone mirror omniview systems can be constructed; this work [7] is an expansion on that important discovery. We show that even under the pinhole camera model the SVP cone configuration works and sees any world point in its FOV, not just rays grazing the mirror surface.

The cone mirror has not been used to construct an SVP omnidirectional imaging sensor that can reproduce perspective projection views from a single omniview image before our work. However, cone mirrors have been used to aid navigation, map building, collision avoidance, and pipe inspection in non-SVP configurations. The cone mirror images were used as is, and no attempt was made to unwarp them to undistorted images. By using multiple normal cameras positioned properly in relation to a plane mirror pyramid, a high-resolution, SVP, wide-FOV system can be built. The trade-offs, though, are the high price and complexities involved with multiple cameras. Bulky size, weight, calibration, synchronization, and gain differences are problems associated with multicamera systems that single-camera systems are free of.

An SVP system is worthwhile if the benefits outweigh the drawbacks for a particular application. Only with SVP can a catadioptric omnicam use a single range-independent lookup table or formula for correct unwarping. The SVP cone system is cheap and simple to build, operate, and maintain while retaining a decent vertical resolution and good flexibility in SVP. The SVP cone system is therefore always worth evaluating before considering more complex and expensive omniview sensors. The main purpose of our work here is to prove that an SVP cone system is both theoretically and physically viable and to present a detailed analysis of cone SVP systems that provides systematic physics-based guidelines for deciding whether the SVP cone is suitable for a particular application.

We analyze here the aberrations of SVP cone mirror systems using accurate numerical optical ray tracing. With our analysis we derive an optical design that minimizes such aberrations. The cone is among the simplest mirror shapes to produce, and it has much higher meridional (tangential) angular resolution compared with other conic section mirrors for scenes around the horizon. It adds the least optical distortion to the resulting meridional images because it is the only omniview mirror with a noncurved mirror surface in the meridional cross sections.

3 Thrust III: Handling Uncertainty and Adversarial Intent in Adaptive Planning.

1. Decentralized nonlinear model predictive control of multiple flying robots.

   We have studied nonlinear model predictive control (NMPC) for multiple autonomous helicopters in a complex environment. NMPC has been recognized as a suitable framework for the control of nonlinear dynamic systems subject to operating contraints. The inherent expandability of NMPC, which can include various performance and constraints as a cost function, has shown promise for the control of unmanned vehicles. However, the computational load of the conventional NMPC was often prohibitive for many dynamic systems with a fast response.

   In our work [13], we proposed to resolve the path planning and optimal control problems for multiple mobile robots in a complex 3-D environment by nonlinear model predictive control and potential function techniques. We have shown that our integrated approach can solve various realistic scenarios involving multiple UAVs and obstacles in a complex 3-D environment.

   The proposed nonlinear trajectory planning and tracking controller has been evaluated in several representative scenarios. Two examples are shown in Figure 4 and 5, respectively.

2. Efficient gradient estimation for motor control learning.

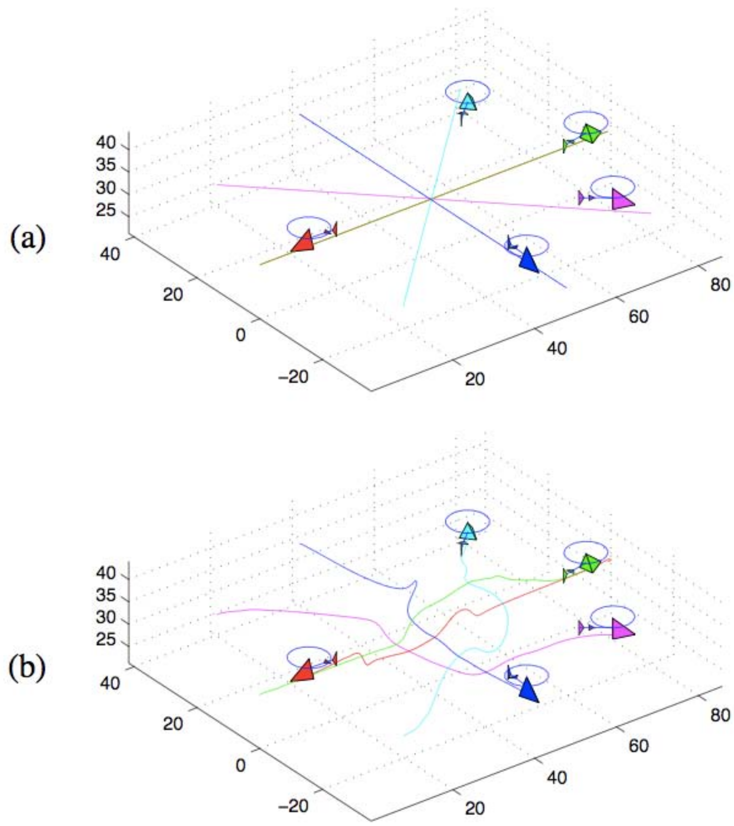

Figure 4: Distributed collision avoidance: (a) initial configuration and the destination of each helicopter, and the corresponding desired trajectories; (b) their trajectories during execution.

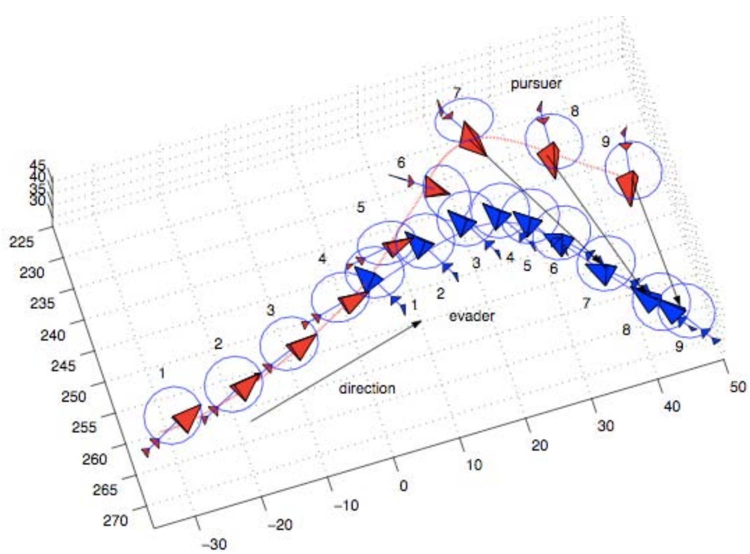

Figure 5: Pursuit-evasion in a 3-D environment: lines between the pursuer (red) and the evader (blue) represent the corresponding positions at each time instant.

The task of estimating the gradient of a function in the presence of noise is central to several forms of reinforcement learning, including policy search methods. We present two techniques for reducing gradient estimation errors in the presence of observable input noise applied to the control signal [6]. the first method extends the idea of a reinforcement baseline by fitting a local model to the response function whose gradient is being estimated; we shown how to find the response surface model that minimizes the variance of the gradient estimation, and how to estimate the model from data. The second method improves this further by discounting components of the gradient vector that have high variance. These methods are applied to the problem of motor control learning, where actuator noise has a significant influence on behavior. In particular, we apply the techniques to learn locally optimal controllers for a dart-throwing task using a simulated three-link arm; we demonstrate that the proposed methods significantly improve the response function gradient estimate and, consequently, the learning curve, over existing methods.

3. Solving models of controlled dynamic planar rigid-body systems with frictional contact.

   The effective operation of robotic systems in the real world critically depends on their reliable interaction with the environment through both intentionally and unintentionally established frictional contacts. Robust manipulation of either the robot or the environment through such interactions, especially for tasks which depend explicitly on accurate control of contact forces, remains a challenging problem for the robotics community. The primary difficulty arises from complexities associated with each of the variety of model choices available for describing interactions between two objects that are in contact.

   In our study, we adopt a class of models that presume rigid objects and Coulomb friction at contact point. It is our conjecture that the resulting models, which still presenting some computational difficulties due to non-linearity and ambiguity, represent an effective and tractable basis for the design of reliable behavioral control systems for contact tasks. The main contribution is a technique and an associated algorithm which computes the instantaneous dynamic response for controlled multi-rigid-body systems subject to Coulomb frictional contacts [3].

   The proposed algorithm makes it possible to predict the range of possible instantaneous body accelerations and contact forces that might result from the application of a specific control input at a specific point in the state space of the system. Our approach is to construct a dynamic response function that is a piecewise linear but possibly ambiguous forward model for the dynamics of the system, which can be locally inverted to aid in the generation of control strategies. In building this function, we also rigorously characterize well-known problems associated with rigid object models involving frictional contacts, such as multiplicity and non-existence of solutions. Consequently, we are able to understand when and how the underlying models fail, allowing controllers to avoid difficult settings or the recognize when they are unavoidable. Specifically, in the presence of possible ambiguities, it becomes possible to design controllers than can either attempt to avoid inputs that result in uncertain behavior, or choose controls that maximize performance in the worse case. The end result is a well-informed control system capable of regulating physical interaction with the environment to produce desired behavior. The multimedia demonstration of this work can be accessed at `http://www.ijrr.org/contents/24_11/abstract/911.html`.

4. Vision-based distributed control laws for motion coordination of nonholonomic robots.

   The main underlying theme of cooperative control of multiple autonomous agents is to analyze and/or synthesize spatially distributed control architectures that can be used for motion coordination of large groups of autonomous vehicles. A key assumption implied in most such applications is that each agent communicates its position and/or velocity information to its neighbors.

   Inspired by the social aggregation phenomena in birds and fish, researchers in robotics and control theory have been developing tools, methods and algorithms for distributed motion coordination of multi-vehicle system. In this project, we propose a set of control laws for coordinated motions such as parallel and circular formations for a group of planar agents using purely local interactions [10]. The derived control laws are in terms of shape variables such as the relative distances and relative headings among the agents. However, these parameters are not readily measurable using simple and basic sensing capabilities. This motivates the rewriting of the derived control laws in terms of biological measurable parameters. Each agent is assumed to have only monocular vision and capable of measuring basic visual quantities such as bearing angle, optical flow, and time-to-collision.

Verification of the theory through multi-robot experiments demonstrated the effectiveness of the vision-based control laws to achieve different formations. We further provide inter-agent collision avoidance properties to the team members. What we show is that the two tasks of formation-keeping and collision-avoidance can be done with decoupled additive terms in the control law, where the terms for keeping parallel and circular formations depend only on visual parameters.

In our experiment, we use a series of small form-factor robots call *Scarab* shown in Figure 6. The Scarab is a $20 \times 13.5 \times 22.2\,cm^3$ indoor ground platform with a mass of 8 kg. Figure 7 shows the trajectories of five robots to form a circular formation with collision avoidance.

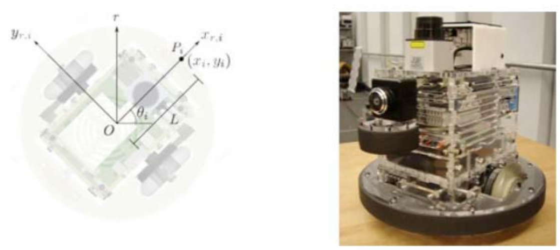

Figure 6: The Scarab is a small robot with a differential drive axle. LED markers are placed on top of each Scarab for pose estimation.

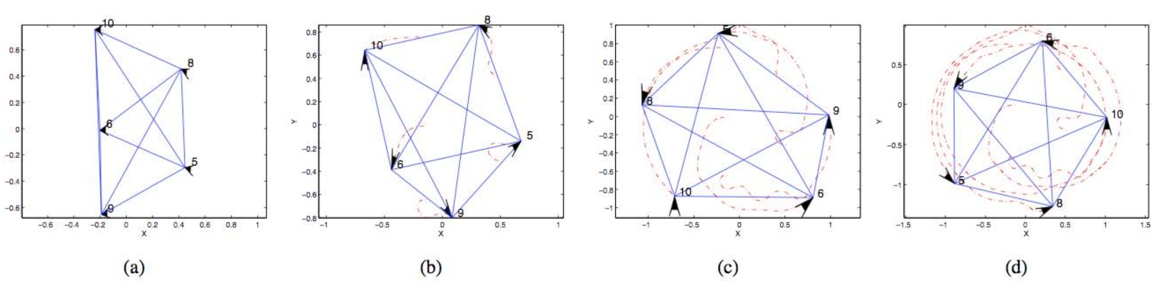

Figure 7: Five Scarabs form a circular formation starting with a complete-graph topology while avoiding collision. (a) $t = 0$ sec. (b) $t = 8$ sec. (c) $t = 20$ sec. (d) At $t = 36$ sec. The robots reach a stable balanced configuration around a circle with radius of 1m. The dash lines show the actual trajectories of the robots and their connectivity graph at the times specified above.

References

[1] A. Antoniades, H. Kim, and S. Sastry, "Pursuit-evasion strategies for teams of multiple agents with incomplete information," in *Proceedings of the IEEE Conference on Decision and Control*, 2003.

[2] J. Barreto and K. Daniilidis, "Wide area multiple camera calibration and estimation of radial distortion," in *The Fifth Workshop on Omnidirectional Vision, Camera Networks and Non-Classical Cameras*, 2004.

[3] A. Greenfield, U. Saranli, and A. Rizzi, "Solving models of controlled dynamic planar rigid-body systems with frictional contact," *International Journal of Robotics Research*, vol. 24, no. 11, 2005.

[4] H. Kim and D. Shim, "A flight control system for aerial robots: algorithms and experiments," *Control Engineering Practice*, vol. 11, no. 12, pp. 1389–1400, 2003.

[5] A. Kumar, J. Tardif, R. Anati, and K. Daniilidis, "Experiments on visual loop closing using vocabulary trees," in *Proceedings of Computer Vision and Pattern Recognition Workshops*, 2008.

[6] G. Lawrence, N. Cowan, and S. Russell, "Efficient gradient estimation for motor control learning," *Uncertainty in Artificial Intelligence*, 2003.

[7] S. Lin and R. Bajcsy, "Single-viewpoint, catadioptric cone mirror omnidirectional imaging theory and analysis," *Journal of Optical Society of America*, vol. 23, no. 12, pp. 2997–3015, 2006.

[8] A. Makadia, A. Patterson, and K. Daniilidis, "Fully automatic registration of 3D point clouds," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2006.

[9] A. Makadia, C. Geyer, and K. Daniilidis, "Correspondence-free structure from motion," *International Journal of Computer Vision*, vol. 75, no. 3, pp. 311–327, 2007.

[10] N. Moshtagh, N. Michael, A. Jadbabaie, and K. Daniilidis, "Vision-based, distributed control laws for motion coordination of nonholonomic robots," *IEEE Transactions of Robotics*, vol. 25, no. 4, pp. 851–860, 2009.

[11] A. Papachristodoulos and A. Jadbabaie, "Synchronization in oscillator networks: Switching topologies and non-homogeneous delays," in *Proceedings of the IEEE Conference on Decision and Control*, 2005.

[12] A. Tahbaz-Salehi and A. Jadbabaie, "On consensus over random networks," *Forty Fourth Annual Allerton Conference*, 2006.

[13] D. Shim, H. Kim, and S. Sastry, "Decentralized nonlinear model predictive control of multiple flying robots," in *Proceedings of the IEEE Conference on Decision and Control*, 2003.

[14] X. Zabulis, A. Patterson, and K. Daniilidis, "Digitizing archaeological excavations from multiple views," in *3-D Digital Imaging and Modeling*, 2005.